

Claims

- [c1] 1. A method for replicating a transaction from a primary database to a replicate database while the replicate database remains available for use, the method comprising:
- recording information about a transaction being performed at a primary database in a transaction log;
- synchronously copying the information about the transaction in the transaction log to a mirrored transaction log;
- generating a reconstructed transaction based on the information about the transaction copied to the mirrored transaction log; and
- applying the reconstructed transaction at the replicate database while the replicate database remains available for use.
- [c2] 2. The method of claim 1, wherein said transaction includes a selected one of a Structured Query Language (SQL) "INSERT", "UPDATE", "DELETE", "DDL" and "Procedure" operation.
- [c3] 3. The method of claim 1, wherein said recording step includes recording at least one log record about the

transaction in the transaction log.

- [c4] 4. The method of claim 3, wherein said at least one log record characterizes changes made to the primary database in the transaction.
- [c5] 5. The method of claim 1, wherein said synchronously copying step includes using a file mirroring module.
- [c6] 6. The method of claim 1, wherein said synchronously copying step includes using file replication hardware.
- [c7] 7. The method of claim 1, wherein said synchronously copying step includes using file replication software.
- [c8] 8. The method of claim 1, wherein said synchronously copying step includes synchronously copying information to the transaction log and the mirrored transaction log before completing the transaction at the primary database.
- [c9] 9. The method of claim 1, wherein said synchronously copying step includes replicating at a file block level the information about the transaction in the transaction log to the mirrored transaction log.
- [c10] 10. The method of claim 1, further comprising:
copying database schema information from the primary database to a site at which the mirrored transaction log

is located to enable transactions to be reconstructed and applied at the replicate database.

- [c11] 11. The method of claim 10, wherein said generating step includes generating the reconstructed transaction based, at least in part, on said database schema information.
- [c12] 12. The method of claim 1, wherein said generating step includes formatting the reconstructed transaction so that the reconstructed transaction is in the same format as the transaction at the primary database.
- [c13] 13. The method of claim 1, wherein said applying step includes verifying that the transaction ordering is correct.
- [c14] 14. The method of claim 1, wherein said applying step includes applying the reconstructed transaction at the replicate database in the same order as the transaction order at the primary database.
- [c15] 15. The method of claim 1, further comprising:
responding to a database query at the replicate database while a transaction is being replicated from the primary database to the replicate database.
- [c16] 16. A computer-readable medium having computer-

executable instructions for performing the method of claim 1.

- [c17] 17. A downloadable set of computer-executable instructions for performing the method of claim 1.
- [c18] 18. A system for replicating transactions from a source database to a standby database, the system comprising:
 - a source database having a transaction log, the transaction log for recording log records for transactions performed at the source database;
 - a mirrored transaction log for recording mirror copies of the log records for transactions performed at the source database;
 - a file mirroring module for synchronously replicating log records from the transaction log to the mirrored transaction log as transactions are performed at the source database;
 - a log reader module for reading log records in the mirrored transaction log and reconstructing transactions for application at the standby database based upon log records in the mirrored transaction log; and
 - a distribution module for applying the transactions reconstructed by the log reader module at the standby database.
- [c19] 19. The system of claim 18, wherein said standby

database is available for responding to database queries while transactions are being replicated from the source database to the standby database.

- [c20] 20. The system of claim 18, wherein said transactions include a selected one of a Structured Query Language (SQL) "INSERT", "UPDATE", "DELETE", "DDL" and "Procedure" operation.
- [c21] 21. The system of claim 18, wherein said log records characterize changes made to the source database based upon transactions performed at the source database.
- [c22] 22. The system of claim 18, wherein said file mirroring module comprises file replication hardware.
- [c23] 23. The system of claim 18, wherein said file mirroring module comprises a disk mirroring module.
- [c24] 24. The system of claim 18, wherein said file mirroring module replicates log records in the transaction log to the mirrored transaction log at a file block level.
- [c25] 25. The system of claim 18, wherein said file mirroring module replicates log records relating to a particular transaction performed at the source database to the mirrored transaction log before said particular transaction is completed at the source database.

- [c26] 26. The system of claim 18, wherein said log reader module reconstructs transactions based, at least in part, on database schema information for the source database.
- [c27] 27. The system of claim 26, further comprising: database schema information for the source database.
- [c28] 28. The system of claim 18, wherein said log reader module formats the reconstructed transactions so that the reconstructed transactions are in the same format as the transaction at the source database.
- [c29] 29. The system of claim 18, wherein said distribution module applies reconstructed transactions at the standby database in the same order as the order of transactions applied at the source database.
- [c30] 30. A method for replicating a database operation from a first database to a second database while making the second database available for decision support purposes, the method comprising:
 - as a database operation is performed at the first database, generating at least one log record characterizing said operation;
 - synchronously recording said at least one log record in a first log associated with the first database and a second

log associated with the first log; and while the second database is available for decision support purposes, replicating said operation performed at the first database at the second database by performing the substeps of: constructing a replicate operation based, at least in part, on said at least one log record in the second log; and applying the replicate operation at the second database.

- [c31] 31. The method of claim 30, wherein said operation includes a selected one of a Structured Query Language (SQL) "INSERT", "UPDATE", "DELETE", "DDL" and "Procedure" operation.
- [c32] 32. The method of claim 30, wherein said synchronously recording step includes file mirroring.
- [c33] 33. The method of claim 30, wherein said synchronously recording step includes replicating said at least one log record to the second log at a file block level.
- [c34] 34. The method of claim 30, wherein said synchronously recording step includes using a disk mirroring module.
- [c35] 35. The method of claim 30, further comprising:
copying database schema information from the first database prior to performing said operation at the first database.

- [c36] 36. The method of claim 35, wherein said constructing substep includes constructing a replicate operation based, at least in part, on said database schema information.
- [c37] 37. The method of claim 35, further comprising:
 - tracking modifications to said database schema information at the first database; and
 - constructing a replicate operation based on said database schema information in effect when the operation is performed at the first database.
- [c38] 38. The method of claim 30, further comprising:
 - assigning a unique identifier to database objects at the first database;
 - if a database object is modified, assigning a different unique identifier to the database object that is modified;
 - and
 - determining a particular database object to be used in constructing a replicate operation based upon said unique identifier assigned to said particular database object.
- [c39] 39. The method of claim 30, wherein said constructing substep includes formatting the replicate operation in the same manner as said operation at the first database.

- [c40] 40. The method of claim 30, wherein said applying sub-step includes applying the replicate operation at the second database in the same order as said operation is applied at the first database.
- [c41] 41. The method of claim 30, wherein making the second database available for decision support purposes includes responding to a database query as said operation is being replicated.
- [c42] 42. The method of claim 30, wherein making the second database available for decision support purposes includes providing access to data in the second database as said operation is being replicated.
- [c43] 43. A method for replicating transactions from a primary database to a replicate database while the replicate database remains available for use, the method comprising:
 - recording log records for transactions being performed at a primary database in a primary transaction log;
 - creating a mirrored transaction log, the mirrored transaction log containing copies of the log records in the primary transaction log;
 - generating reconstructed transactions based on the copies of the log records in the mirrored transaction log;

and applying the reconstructed transactions at the replicate database while the replicate database remains available for use.